

IN THE CLAIMS:

Please amend the claims as follows. The claims are in the format as required by 35 C.F.R. § 1.121.

1. (Currently amended) A system for allocating resources amongst a plurality of applications, the system comprising:

a distributed server pool director for organizing and maintaining a set of servers in a server pool;

a monitoring module at each server in the server pool for detecting demands for one or more resources located on the server and exchanging information regarding demands for the one or more resources at the plurality of servers;

a plurality of computers connected to one another through a network;

a distributed policy engine for specifying a policy for allocation of resources of the plurality of computers amongst [[a]] the plurality of applications having access to the resources, wherein the policy is based on the demands for the one or more resources and adapted based on changes in the behavior of the system;

a monitoring module at each computer for detecting demands for the resources and exchanging information regarding demands for the resources at the plurality of computers; and

an enforcement module at each computer for allocating the resources amongst the plurality of applications based on the policy and information regarding demands for the resources a decision made by the policy engine.

2. (Original) The system of claim 1, wherein the resources include communication resources.

3. (Original) The system of claim 2, wherein the communication resources include network bandwidth.

4. (Original) The system of claim 1, wherein the resources include processing resources.

5. (Original) The system of claim 1, wherein the resources include a selected one of memory, disk space, system I/O (input/output), printers, tape drivers, and software licenses.

6. (Currently amended) The system of claim 1, wherein the policy engine receives user input for defining an application subject to the policy.
7. (Original) The system of claim 6, wherein the monitoring module identifies an application running at a given computer based, at least in part, upon the user input for defining the application.
8. (Original) The system of claim 7, wherein the monitoring module detects a request for resources by the application at the given computer.
9. (Original) The system of claim 6, wherein the user input includes defining components of an application.
10. (Original) The system of claim 9, wherein the components include a selected one of processes, network traffic, and J2EE components.
11. (Original) The system of claim 1, wherein the policy engine receives user input of a policy specifying actions to be taken for allocation of the resources in response to particular conditions.
12. (Original) The system of claim 11, wherein the policy includes a command to be run in response to a particular condition.
13. (Original) The system of claim 11, wherein the policy includes an attribute indicating when a particular condition of the policy is to be evaluated.
14. (Original) The system of claim 13, wherein the policy includes an attribute indicating when action is to be taken based upon a particular condition of the policy being satisfied.
15. (Original) The system of claim 11, wherein the policy specifies priorities of the plurality of applications to the resources.

16. (Original) The system of claim 15, wherein the enforcement module allocates resources amongst the plurality of applications based, at least in part, upon the specified priorities.
17. (Original) The system of claim 1, wherein the policy engine includes a user interface for a user to specify the policy.
18. (Original) The system of claim 1, wherein the policy engine supports an expression language for policy definition.
19. (Original) The system of claim 1, wherein the policy engine is a distributed system operating at each of the plurality of computers.
20. (Original) The system of claim 1, wherein the monitoring module determines resources available at each computer.
21. (Original) The system of claim 1, wherein the monitoring module determines resource utilization at each computer.
22. (Original) The system of claim 21, wherein the monitoring module at each computer exchanges resource utilization information amongst the plurality of computers.
23. (Original) The system of claim 1, wherein the enforcement module allocates network bandwidth amongst said plurality of applications based upon the policy and information regarding demands for the resources.
24. (Original) The system of claim 1, wherein the enforcement module allocates processor resources amongst said plurality of applications based upon the policy and information regarding demands for the resources.
25. (Original) The system of claim 1, wherein the enforcement module includes an interface for communication with an external module for specifying allocation of resources by said external module.

26. (Original) The system of claim 25, wherein said external module includes a load balancer for load balancing instances of an application.

27. (Original) The system of claim 25, wherein said external module comprises a selected one of a router and a provisioning device.

28. (Original) The system of claim 1, wherein the enforcement module starts an instance of an application on a given computer based upon the policy and information regarding demands for the resources.

29. (Withdrawn) An improved method for allocating resources of a plurality of computers to a plurality of applications, the method comprising:  
receiving user input specifying a dynamically configurable policy for allocating resources of a plurality of computers amongst a plurality of applications having access to the resources;  
at each of the plurality of computers, detecting demands for the resources from the plurality of applications and availability of the resources;  
exchanging information regarding demand for the resources and availability of the resources amongst the plurality of computers; and  
allocating the resources to each of the plurality of applications based on the dynamically configurable policy and the information regarding demand for the resources and availability of the resources.

30. (Withdrawn) The method of claim 29, wherein the resources include communication resources.

31. (Withdrawn) The method of claim 30, wherein the communication resources include network bandwidth.

32. (Withdrawn) The method of claim 29, wherein the resources include processing resources.

33. (Withdrawn) The method of claim 29, wherein the resources include a selected one of memory, disk space, system I/O (input/output), printers, tape drivers, load balancers, and software licenses.

34. ((Withdrawn) The method of claim 29, wherein said receiving step includes receiving user input for defining an application.

35. (Withdrawn) The method of claim 34, further comprising:  
providing a set of default rules for assisting a user in defining an application.

36. (Withdrawn) The method of claim 29, wherein said detecting step includes detecting applications running on each of said plurality of computers.

37. (Withdrawn) The method of claim 29, wherein said detecting step associating particular processes with a particular application.

38. (Withdrawn) The method of claim 29, wherein said detecting step associating particular network traffic with a particular application.

39. (Withdrawn) The method of claim 29, wherein said detecting step includes detecting components of an application.

40. (Withdrawn) The method of claim 39, wherein said step of detecting components of an application includes detecting a J2EE component.

41. (Withdrawn) The method of claim 29, wherein said receiving step includes receiving user input specifying actions to be taken for allocation of the resources in response to particular conditions.

42. (Withdrawn) The method of claim 41, wherein the user specifies a script to be run based upon a particular condition.

43. (Withdrawn) The method of claim 41, wherein the user specifies when a particular condition of the dynamically configurable policy is to be evaluated.
44. (Withdrawn) The method of claim 43, wherein the user specifies when action is to be taken based upon the particular condition being satisfied.
45. (Withdrawn) The method of claim 29, wherein said receiving step includes receiving user input specifying priorities of the plurality of applications to the resources.
46. (Withdrawn) The method of claim 45, wherein the allocating step includes allocating resources amongst the plurality of applications based, at least in part, upon the specified priorities.
47. (Withdrawn) The method of claim 29, wherein the receiving step includes providing an expression language for policy definition.
48. (Withdrawn) The method of claim 29, wherein said detecting step includes determining resource utilization at the given computer.
49. (Withdrawn) The method of claim 48, wherein said determining resource utilization step includes determining average resource utilization over a given time period.
50. (Withdrawn) The method of claim 29, wherein said allocating step includes allocating a specified amount of resources to a particular application when the particular application is initially detected at a given computer.
51. (Withdrawn) The method of claim 29, wherein said allocating step includes allocating resources based upon particular events.
52. (Withdrawn) The method of claim 29, wherein said allocating step includes allocating network bandwidth to each of the applications.

53. (Withdrawn) The method of claim 29, wherein said allocating step includes allocating processor resources amongst the plurality of applications.

54. (Withdrawn) The method of claim 29, wherein said allocating step includes communicating with an external module for allocating resources provided by an external module.

55. (Withdrawn) The method of claim 54, wherein the external module comprises a load balancer for load balancing in instances of an application.

56. (Withdrawn) The method of claim 54, wherein the external module comprises a selected one of a router and a provisioning device.

57. (Withdrawn) The method of claim 29, wherein said allocating step includes starting an instance of an application on a given computer.

58. (Withdrawn) A computer-readable medium having processor executable instructions for performing the method of claim 29.

59. (Withdrawn) A downloadable set of processor-executable instructions for performing the method of claim 29.

60. (Currently amended) A method for allocating resources to a plurality of applications, the method comprising:

receiving user input specifying priorities of the plurality of applications to resources of a plurality of servers, the specified priorities including designated servers assigned to at least some of the plurality of applications;

selecting a given application based upon the specified priorities of the plurality of applications;

determining the demand for one or more resources located on each server in a plurality of communicatively connected servers;

specifying a policy for allocation of resources of the plurality of servers;

~~determining available servers on which the given application is runnable and which are not assigned to a higher priority application;~~

~~allocating to the given application any available resources located on one or more of the plurality of servers which are designated servers assigned to the given application based on the policy for allocation of resources;~~

~~allocating any additional available servers resources to the given application until the given application's demands for resources are satisfied; and~~

~~repeating above steps for each of the plurality of applications based on the specified priorities.~~

61. (Original) The method claim 60, wherein the receiving step includes receiving user input of a value for a given application representing relative priority of the given application compared to other applications.

62. (Original) The method of claim 60, wherein the receiving step includes receiving a ranking of the plurality of applications from highest priority to lowest priority.

63. (Original) The method of claim 62, wherein the step of selecting a given application includes commencing with selection of an application having the highest priority.

64. (Original) The method of claim 60, further comprising:  
powering on a server allocated to an application if the server is in a powered off state.

65. (Original) The method of claim 60, further comprising:  
determining whether an application is inactive on a server allocated to the application;  
and  
initiating a resume script for running the application on the server application is determined to be inactive.

66. (Original) The method of claim 60, further comprising:  
adding a server newly allocated to an application to a set of servers across which the application is load balanced.

67. (Currently amended) The method of claim 60, further comprising:  
removing a ~~sewer~~ server no longer allocated to an application from a set of servers across which the application is load balanced.
68. (Original) The method of claim 60, further comprising:  
determining whether a server no longer allocated to an application is in a suspend set of servers designated for the application; and  
running a suspend script if the server is determined to be in the suspend set of servers.
69. (Original) The method of claim 68, further comprising:  
if a suspend script is executed on the server, determining whether the server should be powered off based on consulting a power management rule; and  
powering off the server if it determined that the server should be powered off.
70. (Original) The method of claim 60, wherein said allocating step includes starting an instance of an application on a given computer.
71. (Original) A computer-readable medium having processor executable instructions for performing the method of claim 60.
72. (Cancelled).